Clouds and the Earth's Radiant Energy System (CERES)

Data Management System

Software Design Document

Tisa Averaging

(Subsystems 7.1, 8.0, and 10.0)

Architectural Draft

Kay Costulis[1], Maria Mitchum[1], Raja Raju[2],
Libby Smith[2], Sarah Sullivan[2], and Jill Travers[2]

[1]NASA Langley Research Center
Mail Stop 423
Hampton, VA 23681-0001

[2]Science Applications International Corporation (SAIC)
One Enterprise Parkway, Suite 300
Hampton, Virginia 23666

Atmospheric Sciences Division
NASA Langley Research Center
Hampton, Virginia 23681-0001

Release 1.0
June 1996

Preface

The Clouds and the Earth's Radiant Energy System (CERES) Data Management System supports the data processing needs of the CERES science research to increase understanding of the Earth's climate and radiant environment. The CERES Data Management Team works with the CERES Science Team to develop the software necessary to support the science algorithms. This software, being developed to operate at the Langley Distributed Active Archive Center (DAAC), produces an extensive set of science data products.

The Data Management System consists of 12 subsystems; each subsystem represents a stand-alone executable program. Each subsystem executes when all of its required input data sets are available and produces one or more archival science products.

The documentation for each subsystem describes the software design at various stages of the development process and includes items such as Software Requirements Documents, Data Products Catalogs, Software Design Documents, Software Test Plans, and User's Guides.

This version of the Software Design Document records the architectural design of each Subsystem for Release 1 code development and testing of the CERES science algorithms. This is a PRELIMINARY document, intended for internal distribution only. Its primary purpose is to record what was done to accomplish Release 1 development and to be used as a reference for Release 2 development.

Acknowledgements are made to David F. Young and David R. Doelling, Analytical Services & Materials, Inc., for their science support.

TABLE OF CONTENTS

TABLE OF CONTENTS

Section | Page

List of Tables

List of Figures

1.0 Introduction

The Clouds and the Earth's Radiant Energy System (CERES) is a key component of the Earth Observing System (EOS). The CERES instruments are improved models of the Earth Radiation Budget Experiment (ERBE) scanner instruments, which operated from 1984 through 1990 on the National Aeronautics and Space Administration's (NASA) Earth Radiation Budget Satellite (ERBS) and on the National Oceanic and Atmospheric Administration's (NOAA) operational weather satellites NOAA-9 and NOAA-10. The strategy of flying instruments on Sun-synchronous, polar orbiting satellites, such as NOAA-9 and NOAA-10, simultaneously with instruments on satellites that have precessing orbits in lower inclinations, such as the Earth Radiation Budget Satellite (ERBS), was successfully developed in ERBE to reduce time sampling errors. CERES will continue that strategy by flying instruments on the polar orbiting EOS platforms simultaneously with an instrument on the Tropical Rainfall Measuring Mission (TRMM) spacecraft, which has an orbital inclination of 35 degrees. In addition, to reduce the uncertainty in data interpretation and to improve the consistency between the cloud parameters and the radiation fields, CERES will include cloud imager data and other atmospheric parameters. The first CERES instrument is scheduled to be launched on the TRMM spacecraft in 1997. Additional CERES instruments will fly on the EOS-AM platforms, the first of which is scheduled for launch in 1998, and on the EOS-PM platforms, the first of which is scheduled for launch in 2000.

1.1 Document Overview

The purpose of this document is to provide a basis and an explanation for the design of time interpolation processes and the monthly regional, zonal, and global averaging processes for the following three Subsystems: Time and Interpolation for Single and Multiple Satellites (7.1); Compute Regional, Zonal, and Global Averages (8.0); and Compute Monthly and Regional Top-of-the-Atmosphere (TOA) and Surface Radiation Budget (SRB) Averages Subsystem (10.0). The Release 1 design document addresses the requirements from the CERES Science Team's Algorithm Theoretical Basis Documents (References 1 - 4) and the Software Requirements Document for Subsystem 7.1 (Reference 5), Subsystem 8.0 (Reference 6), and Subsystem 10.0 (Reference 7). The intended audience for this document consists of the following groups: the Subsystem design team, programmers, testers, follow-on subsystems, and science reviewers. It serves as a reference so that users may quickly determine the Subsystem's purpose and structure.

This document contains the following information:

- A brief overview of purpose and functionality of these Subsystems.
- A description of key concepts embodied in these Subsystems.
- A description of any constraints on the design and implementation of the software.
- A description of the design approach for these Subsystems.
- The architectural design which includes flow charts, context diagrams, and scenario diagrams describing the structure and flow of these Subsystems.
- References

- Abbreviations and Acronyms (Appendix A).

This document will be updated after Release 1 to reflect changes in the implementation process.

1.2 Subsystem Overviews

1.2.1 Time Interpolation for Single and Multiple Satellites (Subsystem 7.1) Overview

The time interpolation process (7.1), one of the two key parts of Subsystem 7.0, temporally interpolates CERES data and produces global synoptic maps of Top-of-the-Atmosphere (TOA) fluxes and cloud properties on a 1.25-degree equal-area grid. Another key part of Subsystem 7.0, the Synoptic Surface and Atmospheric Radiation Budget (SARB) (Subsystem 7.2), calculates synoptic maps of the vertical structure of atmospheric and surface flux using the interpolated data as input and boundary conditions. These synoptic maps are stored in the product, Synoptic Radiative Fluxes and Clouds (SYN).

The main input to the time interpolation process is the Hourly Gridded Single Satellite Fluxes and Clouds (FSW) product, produced by Grid Single Satellite Radiative Fluxes and Clouds (Subsystem 6.0). The gridded shortwave (SW) and longwave (LW) TOA fluxes and cloud information are the key items to be interpolated. This process produces the internal product, Time Space Interpolate (TSI), which is the input to Subsystem 7.2. The radiative profile will be recalculated in the SARB part of Subsystem 7.0 using the interpolated fluxes as constraints.

The time interpolation process produces global maps of TOA total-sky LW and SW flux, TOA clear-sky LW and SW flux, TOA window radiances, and cloud properties at 0, 3, 6, ..., 21 Greenwich mean time (GMT) for every day of the month. The process of producing synoptic maps involves several steps:

1. Read the FSW data.
2. Read the gridded geostationary data from the Grid ISCCP Geostationary Radiances (GGEO) product.
3. Cloud properties from the CERES times of observation are interpolated to the synoptic times.
4. The CERES TOA LW and SW fluxes are interpolated to synoptic times using geostationary data to assist in modeling meteorological variations between times of observations.

1.2.2 Compute Regional, Zonal, and Global Averages (Subsystem 8.0) Overview

The Compute Regional, Zonal, and Global Averages Subsystem produces regional, zonal and global monthly and monthly-hourly means. These means are calculated from one month of synoptic maps on a regional basis and then combined to produce zonal and global averages. The main input to this Subsystem is the SYN product produced by Subsystem 7.2. This product contains one month of 3-hourly synoptic maps of TOA LW and SW fluxes (total-sky and clear-sky), TOA window radiances, upwelling and downwelling SW and LW flux at each standard CERES pressure level, and numerous cloud parameters for each region of the CERES global 1.25-degree equal-area grid.

The two archival products output from this Subsystem are the Monthly Regional Radiative Fluxes and Clouds (AVG) product which contains regional monthly and monthly-hourly means of fluxes and cloud parameters and the Monthly Zonal and Global Radiative Fluxes and Clouds (ZAVG) product which contains the zonal and global monthly and monthly-hourly averages of the above parameters.

The main steps involved in the averaging process are

1. Read the synoptically ordered data.
2. Average the flux data to produce regional monthly and monthly-hourly means.
3. Average the cloud properties using the specified weighting schemes to produce regional monthly and monthly-hourly means.
4. Average the regional means to produce zonal means.
5. Average the zonal means to produce global means.

1.2.3 Compute Monthly and Regional TOA and SRB Averages (Subsystem 10.0) Overview

The Compute Monthly and Regional TOA and SRB Averages (Subsystem 10.0) computes averages of TOA LW and SW fluxes, surface fluxes, and cloud properties on regional, zonal, and global spatial scales. The main input to Subsystem 10.0 is the SFC product produced by Grid TOA and Surface Fluxes (Subsystem 9.0). SFC contains hourly single satellite flux and cloud properties averaged over 1.25-degree regions. This Subsystem produces the Monthly Regional TOA and SRB Averages (SRBAVG) product. Two methods are used to compute the regional TOA total-sky flux averages. TOA flux estimates from both of the two methods are used to produce estimates of surface flux at all temporal and spatial scales using the TOA-to-surface flux parameterization schemes for shortwave and longwave described in the ERBE-like Inversion to Instantaneous TOA Fluxes (Subsystem 2.0) (Reference 8).

The process of producing the means stored in SRBAVG involves several steps:

1. The ancillary geostationary data which are used in the interpolation of TOA fluxes are gridded to the CERES grid system and regionally and temporally sorted and merged.

2. The TOA total-sky flux data are interpolated by two methods: the ERBE-like method and the geostationary data enhancement method (Reference 9).
3. The TOA clear-sky flux data, surface flux data, and the cloud property data are linearly interpolated.
4. Monthly and monthly hourly means are calculated from the interpolated fluxes and cloud properties on regional, zonal, and global scales.

1.3 Key Concepts

CERES 1.25-degree Equal-area Grid
Input Data Organization
Interpolation
Monthly Hourly and Monthly Means
Regional Averaging
Zonal and Global Averaging
Output Data Organization
Geostationary Data
Solar Information
Area Weighting Factors
Metadata and Error Handling

CERES 1.25-degree Equal-area Grid. "The 1.25-deg equal-area reference grid is constructed by dividing the latitude into 1.25-deg zones and then dividing each zone into an integer number of longitude increments so that the reference regions have nearly equal area. The reference regions at the equator are defined by 1.25-deg increments in both latitude and longitude. Other zones are divided into an integer number of regions that produce approximately the same area as at the equator. Table 1-1 gives the number of reference regions in each zone from the south pole to the equator .

Table 1-1. Number of Reference Grid Regions in Each Zone

South - 3	10 - 59	19 - 113	28 - 163	37 - 206	46 - 241	55 - 267	64 - 283
2 - 9	11 - 65	20 - 119	29 - 168	38 - 210	47 - 245	56 - 270	65 - 284
3 - 16	12 - 72	21 - 125	30 - 173	39 - 214	48 - 248	57 - 272	66 - 285
4 - 22	13 - 78	22 - 130	31 - 178	40 - 219	49 - 251	58 - 274	67 - 286
5 - 28	14 - 84	23 - 136	32 - 183	41 - 223	50 - 254	59 - 276	68 - 287
6 - 34	15 - 90	24 - 141	33 - 188	42 - 227	51 - 257	60 - 277	69 - 287
7 - 41	16 - 96	25 - 147	34 - 192	43 - 230	52 - 260	61 - 279	70 - 288
8 - 47	17 - 101	26 - 152	35 - 197	44 - 234	53 - 262	62 - 280	71 - 288
9 - 53	18 - 107	27 - 157	36 - 201	45 - 238	54 - 265	63 - 282	Eq. - 288

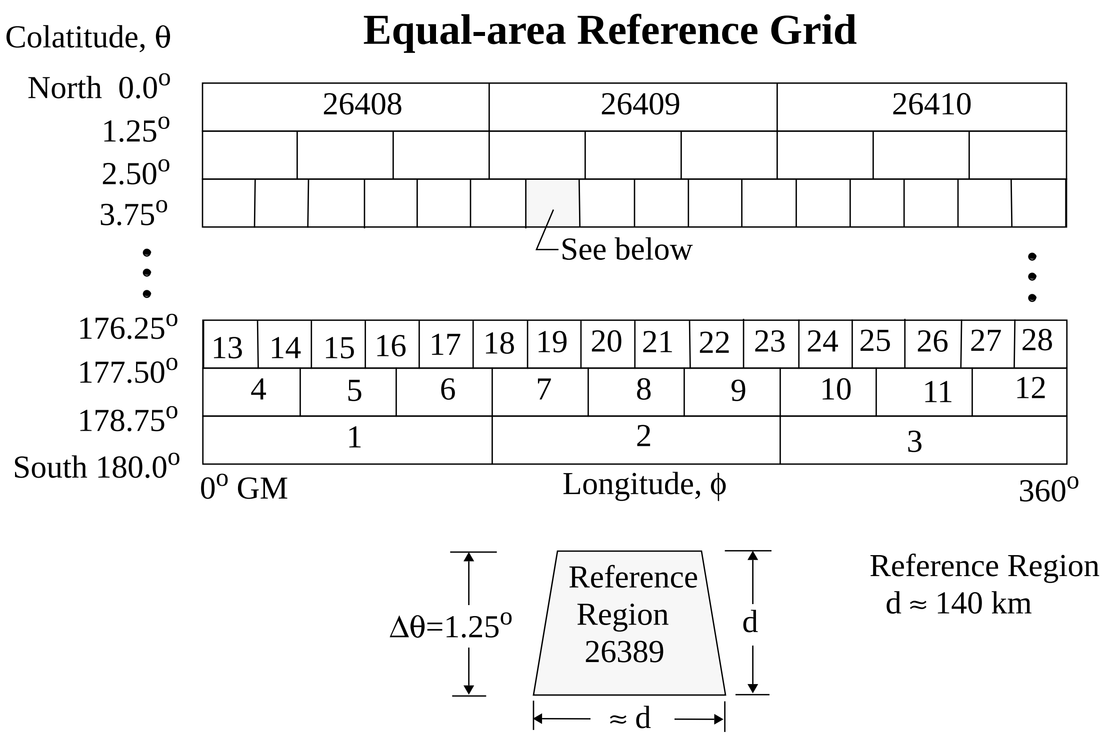

The Reference Grid is symmetric about the equator and contains a total of 26,410 regions. By convention each reference region contains its southern and western boundary. The one exception is at the north pole, where regions contain both their northern and southern boundaries. We start at the south pole and divide the colatitude zone from 180-deg to 178.75-deg into three regions. The first region is at the south pole and has the Greenwich meridian as its western boundary. Successive regions are numbered eastward and northward. This construction preserves the Greenwich meridian and the Earth equator as regional boundaries, but it does not preserve the international date line. Another convention that will be helpful is to define the colatitude and longitude of a region at the "center" of the region" (Reference 10).

Input Data Organization.

- The primary input for Subsystem 7.1, FSW, consists of 1-hour averages of CERES measurements for each observed region of the 1.25-degree equal-area grid. Data for one month (31 days x 24 hours) have been written to 144 zonal files by Grid Single Satellite Radiative Fluxes and Clouds (Subsystem 6.0). These data have been sorted by region and time.
- The primary input for Subsystem 8.0, SYN, consists of the interpolated synoptic maps of the vertical structure of top-of-atmoshpere and surface fluxes, surface parameters, and cloud properties for each region of the 1.25-degree equal-area grid. Data for one month (31 days x 8 synoptic hours) have been written to day files, as well as a secondary index file that stores the record numbers for all the hours of every region. (Refer to Appendix B for more information on these files.)

• The primary input for Subsystem 10.0, SFC, consists of data for all the observed hours for each region for an entire month. SFC is stored as multiple SFC records, where each record contains one hour's worth of measurements for all of the parameters for a particular region. These SFC records are organized in 144 zonal files, where all the hours for all the regions of a zone are stored in one zonal file.

Interpolation. The interpolation process is used in Subsystems 7.1 and 10. It is performed on a regional basis and involves two steps. The first step involves reading in a region's worth of data from the input product. Each parameter from the input product is defined as an array of length 744 for all of the hours in the month. Data from the input products are accessed on an hourly basis for each parameter. Access routines allow Subsystems 7.1 or 10 to read in a parameter's data for each hour a measurement was taken. Each piece of data is placed in the appropriate array where the array index corresponds to the hour box number. This is done for all parameters for the entire month for a region. When all of the data for a region over a month has been read, each parameter has an array containing data only for those hours where data was measured. The second step in the interpolation process fills in the empty hours in the array using various interpolation algorithms, which are dependent on the type of data.

Monthly and Monthly Hourly Means. A monthly hourly mean is the average of all the days in the month for each hour exclusively. For example, the monthly hourly mean for hour 1 is the average of all of the first hours in every day of the month. The grand monthly mean is the average of the monthly hourly means. Figure 1-1 depicts the calculation of these means; an "x" indicates where a measurement occurred, and a bold asterisk(*****) indicates where a monthly hourly mean was calculated.

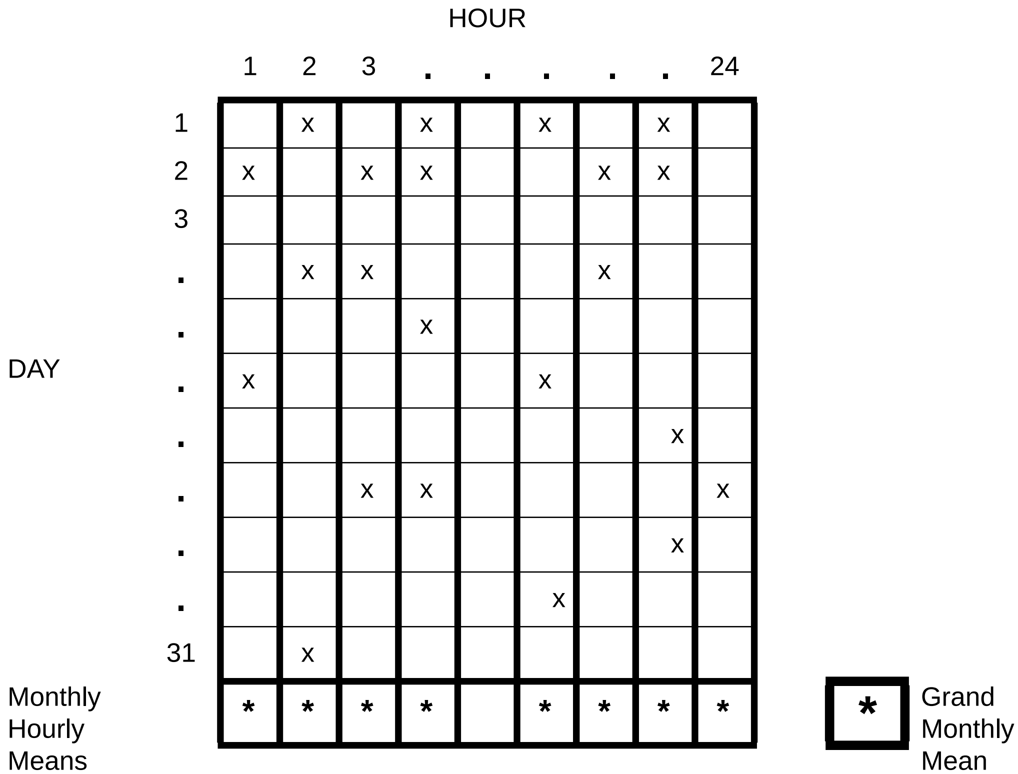

Figure 1-1. Region Data Over a Month for One Parameter

Regional Averaging. Regional averaging is driven by an averaging subroutine that processes a month's worth of data for one region at a time. For every parameter in the region, the averaging routine calls the appropriate averaging algorithm, which depends on the type of data. The averaging algorithm takes all of the interpolated data over the month for the parameter and produces a monthly hourly and monthly mean for that parameter.

Zonal and Global Averaging. The monthly and monthly hourly regional means are further averaged zonally and globally. The zonal average is the average of all regions within a given latitude band or zone. The global average is the average of all of the zones. Area weighting factors are used in the calculation of the global average to correct for the slight variation of grid box size with latitude.

Output Data Organization.

- The output of Subsystem 7.1 is the TSI product, which is organized as 40 files. Each file contains data records for a specific portion of the globe. The synoptic data records are produced at 3-hour intervals in GMT. Therefore, there are a maximum 248 synoptic data records per region for an entire month of data. Data records are written to TSI for only those regions which are observed. The TSI product also has an additional secondary index file that contains the appropriate record number on one of the 40 files for a particular region and synoptic hour. (Refer to Appendix C for more information on these files.)

- The output of Subsystem 8.0 is the AVG/ZAVG product. Regional means are computed for all of the synoptic hours in the month (1..248). Once regional means are computed for all parameters and all regions, they are written to the output product AVG. The regional means are then combined into zonal and global means. These means are then written to the second archival product, ZAVG.
- The output of Subsystem 10 is the SRBAVG product. One SRBAVG record consists of the monthly hourly and monthly averages for every parameter. The averages for one SRBAVG record can be for a region, zone, or the globe. The entire SRBAVG product consists of averages for all of the regions (26,410 SRBAVG records), averages for all of the zones (144 SRBAVG records), and averages for the globe (1 SRBAVG record). The SRBAVG product is currently written to one file.

Geostationary Data. Geostationary data will be the International Satellite Cloud Climatology Project (ISCCP) B3 data for Release 1. The ISCCP B1 data will be used for Release 2. This data is provided at 3-hour intervals at 0,3,6, ... ,21 GMT for a month. The data will be sorted and averaged onto the CERES 1.25-degree equal-area grid and interpolated to local time. This data will be used for interpolating the TOA total-sky longwave and shortwave fluxes.

Solar Information. The solar declination angle and the distance correction constants for each day of the year are produced from a subroutine in CERES library (CERESlib), which is called when the Subsystem is being initialized. This solar information is stored in the sol_stats modules and used to calculate other solar statistics according to the zone being processed.

Area Weighting Factors. The area weighting factors are calculated and obtained for each zone from a subroutine in CERESlib. The area weighting factors are used to weight the means when calculating global averages to correct for the slight variation of grid box size with latitude.

Metadata and Error Handling. Error handling is implemented using the Science Data Production (SDP) Toolkit provided by Earth Observing System Data and Information System (EOSDIS) (Reference 11). Access to the toolkit in the code is through FORTRAN90 Toolkit Wrappers which reside in CERESlib. Metadata is scheduled to be implemented in Release 2.

1.4 Implementation Constraints

Use of the EOSDIS Planning and Data Production System (PDPS) operating environment, the SDP Toolkit, the limits of computer Control Processing Unit (CPU), data throughput, network capacity, and system complexity provide constraints on the design.

Design of the input products FSW and SFC, for Subsystems 7.1 and 10.0, respectively, forces the code for these Subsystems to read in one data measurement for one hour at a time.

At least one month of data is needed in order to process, since Subsystems 8 and 10 produce monthly means, and Subsystem 7.1 requires data for a month for the interpolation algorithms.

Although the archived output products, SRBAVG for Subsystem 10 and AVG/ZAVG for Subsystem 8, are not currently implemented in Hierarchical Data Format (HDF), developers considered the impact HDF implementation would have on this data. The design of the products' structures were constrained by the size limit of HDF structures; for this reason the developers were not able to group the monthly hourly and monthly means together. The design of the output products reflect these constraints.

1.5 Design Approach

In approaching the design for Subsystems 7.1, 8.0, and 10.0, it became apparent that there were enough similarities between these Subsystems that it was possible to design one system that would execute all three Subsystems. First, the similarities and differences between the three subsystems will be discussed. Then we will address FORTRAN90, the language used to implement this system, and the capabilities it provided. We will discuss the design approach of these Subsystems, which is separated into two sections, the data design and the algorithm design, and lastly, a brief discussion on Input/Output (IO) design.

Similarities between Subsystems 7.1, 8.0, and 10.0.
There are many similarities between Subsystems 7.1, 8.0, and 10.0. The similarities can be viewed in terms of functionality and data. The function of these three Subsystems involves the following steps: read in the input product, process every zone on the globe, process every region in every zone, interpolate the missing hours for a region over a month of data, average the interpolated data for a region to produce monthly and monthly hourly means, average the regional/zonal averages to produce zonal/global monthly and monthly hourly means, and write to the output product. First, Subsystem 10.0 executes all of these steps. However, Subsystem 7.1 and Subsystem 8.0 are really parts of 10.0; Subsystem 7.1 performs the interpolation part, and Subsystem 8.0 performs the averaging part. One can see the same functionality in these Subsystems. We then examined the differences between the Subsystems; a discussion follows.

In the CERES Top-level Data Flow Diagram, it can be seen that Subsystems 7.1 and 10.0 have two different input products. The differences between these input products, FSW and SFC, is that the data in FSW are in GMT and the data in SFC are in local time. Also, Subsystems 7.1 and 8.0 were designed to produce synoptic (data every 3 hours in a day) output products, while Subsystem 10.0 was designed to produce an output product that contains means for every hour in the day. Subsystem 8.0 also contains additional parameters that are calculated in Subsystem 7.2 (SARB) and need to be averaged and written to the output product for 8.0, AVG/ZAVG. However, despite these differences the majority of the algorithms to interpolate and average the data are the same. Likewise, all three Subsystems also contain much of the same data, which will be discussed next.

The data that is common to the three Subsystems are location data, top-of-atmosphere fluxes, angular model scene type data, surface only data, and weighted column averaged cloud data. There is also data that is common to Subsystems 7.1 and 8.0: cloud data at the four layers and the cloud overlap conditions. An abstract data type was built for each data type mentioned above. Regardless of which subsystem is running, the abstract data type can be used to declare an instance of the data type, interpolate that data, or average that data. These abstract data types will be discussed in more detail in the following sections.

Use of Fortran90. All data and subroutines for this system are implemented using FORTRAN90 (F90) modules. The F90 module provides the flexiblity to implement several important programming concepts. The first concept is the abstract data type. Using a F90 module, each group of data (TOA fluxes, cloud data, etc.) were implemented as an abstract data type, which will be discussed next under data design. Another programming feature that can be implemented using a F90 module is that of a library of common routines. The F90 module also provides a convenient way to implement science algorithms, whether for interpolation or averaging, where the complexity of the algorithm is hidden to the user.

Data design. FORTRAN90 modules provide a mechanism for building abstract data types. Subsystems 7.1, 8.0, and 10.0 contain many groups of data mentioned earlier. Each data group is built as an abstract data type that can be used by any one of the three subsystems. An abstract data type for a data group consists of the type definitions that defines the structure of the data group, default parameters for the data, a routine to interpolate the missing hours for each parameter in the data group, a subroutine that calls the algorithms to calculate regional means for each parameter in the data group, and a subroutine that calls the algorithms to calculate the zonal or global means for each parameter in the data group. A module implemented as an abstract data type allows the user to declare instances of that data type to be used throughout their code and perform either interpolation, regional averaging, or zonal/global averaging on that data. For example, once the data module for the top-of-atmosphere fluxes has been built, Subsystem 10.0 can declare an instance of TOA fluxes, interpolate that over the month, and average the interpolated data. Likewise, Subsystem 7.1 can use the same data module to declare an instance of TOA fluxes and interpolate that data over the month. Lastly, Subsystem 8.0 can use the same code to average the TOA flux data.

Algorithm design. The FORTRAN90 module also provides a convenient way to organize subroutines. The design of this system contains driver modules, library modules, and complex algorithm modules. Driver modules consist of a main public subroutine that calls private subroutines within the module, as well as subroutines from other modules. A library module contains subroutines that perform similar functions, which are used throughout the code. Lastly, a complex algorithm module provides a public interface subroutine to the algorithm and contains all of the private supporting routines needed to execute the algorithm.

Input/Output (IO) design. One other important aspect of the design of this system is the way that the input products and output products are read and written. There are two modules, one for reading the input data and one for writing the output product. These modules provide generic routines so that the calling program can request data, without knowing from which file it is being read and write out data without knowing to which file it is being written. The system is built so that the IO modules know which Subsystem is running and can retrieve and store data to the appropriate products.

After analyzing the many similarities between Subsystems 7.1, 8.0, and 10.0, one system was designed for all three Subsystems. An object-oriented design provided the concepts to build one system that would be able to execute any of the three Subsystems. It was possible to define the structure and behavior of the data in a data module, so that any subsystem could use the types and subroutines to manipulate the data. The use of IO modules and generic routines precludes the driver subroutines from having to know which files need to be opened for reading or writing. Being able to maintain one set of code for three subsystems will prove to be beneficial in the future.

2.0 Architectural Design

This section gives a top level view of the architectural design and functionality of Subsystems 7.1, 8.0, and 10.0. A flow diagram is given for each subsystem to show which functions are performed for each subsystem. A context diagram is given for each subsystem to show which input and output product modules are used for each subsystem, as well as to demonstrate which data modules are used in the individual subsystems. In Section 2.4, scenario diagrams are given to demonstrate what is processed at the many levels of this system. Since the scenario diagrams for Subsystems 7.1 and 8.0 are really sub-diagrams of the scenario diagram for Subsystem 10.0, only one scenario diagram is given for each level to describe all three of the subsystems. The text describing each step in the scenario diagram specifies which subsystem would execute that call.

2.1 Subsystem 7.1

The main processing flow of Subsystem 7.1 is shown below in the flow diagram in Figure 2-1. This diagram gives an overview of the three main loops and the functions performed within these loops.

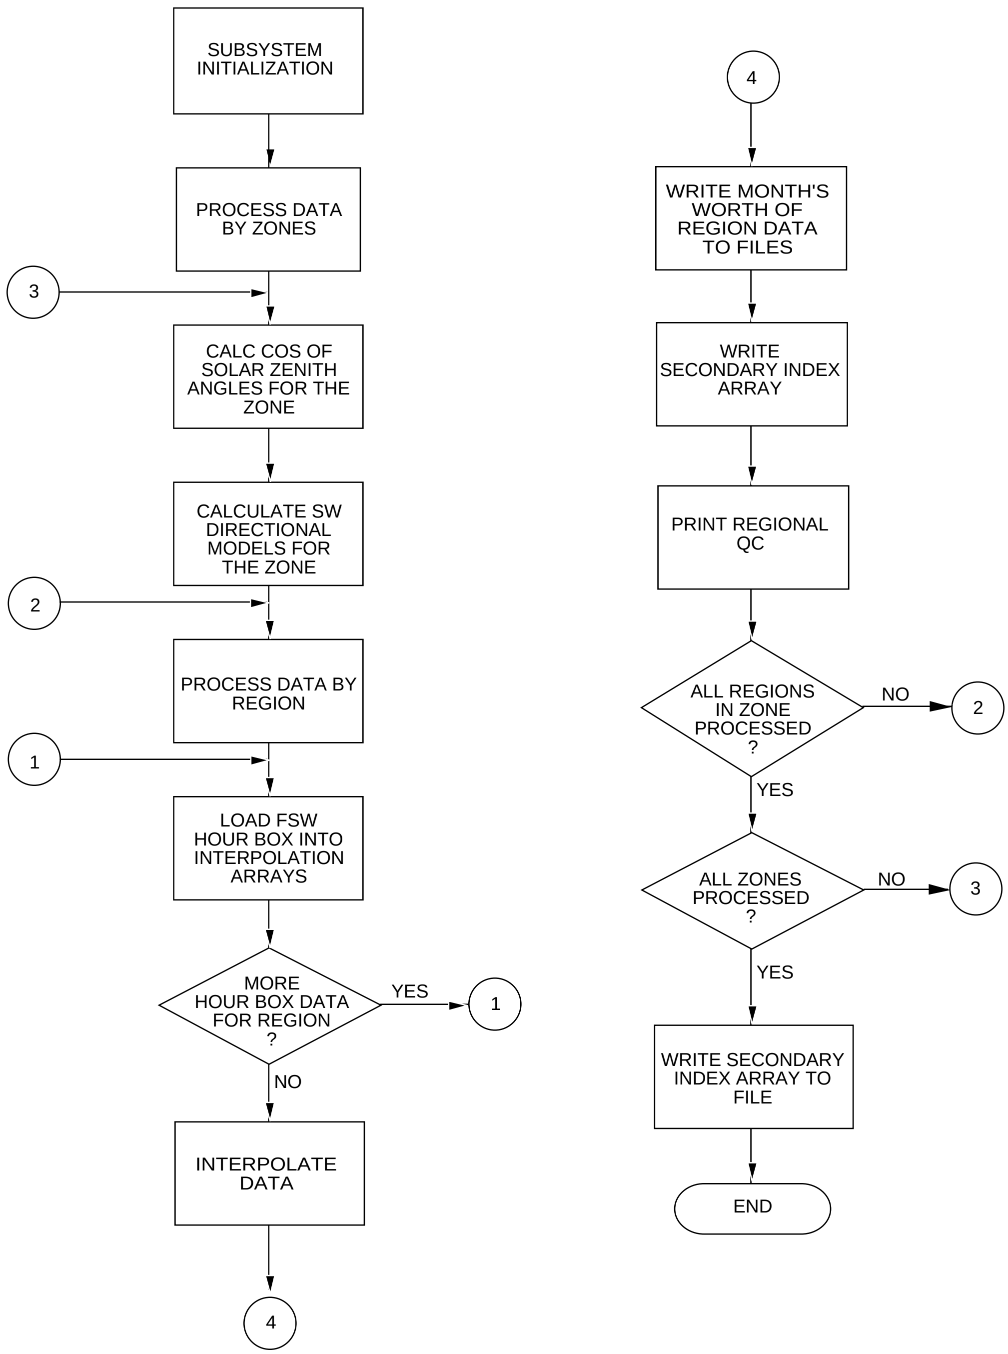

Figure 2-1. Subsystem 7.1 Flow Chart Overview

To introduce the main components of Subsystem 7.1, a context diagram is shown in Figure 2-2. This diagram demonstrates the dependencies between the top-level modules that define the data used in Subsystem 7.1 and control the main processing, which is as follows: initialize the Subsystem; read in the input data, FSW; interpolate the data; and write the synoptic averages to the output product, TSI.

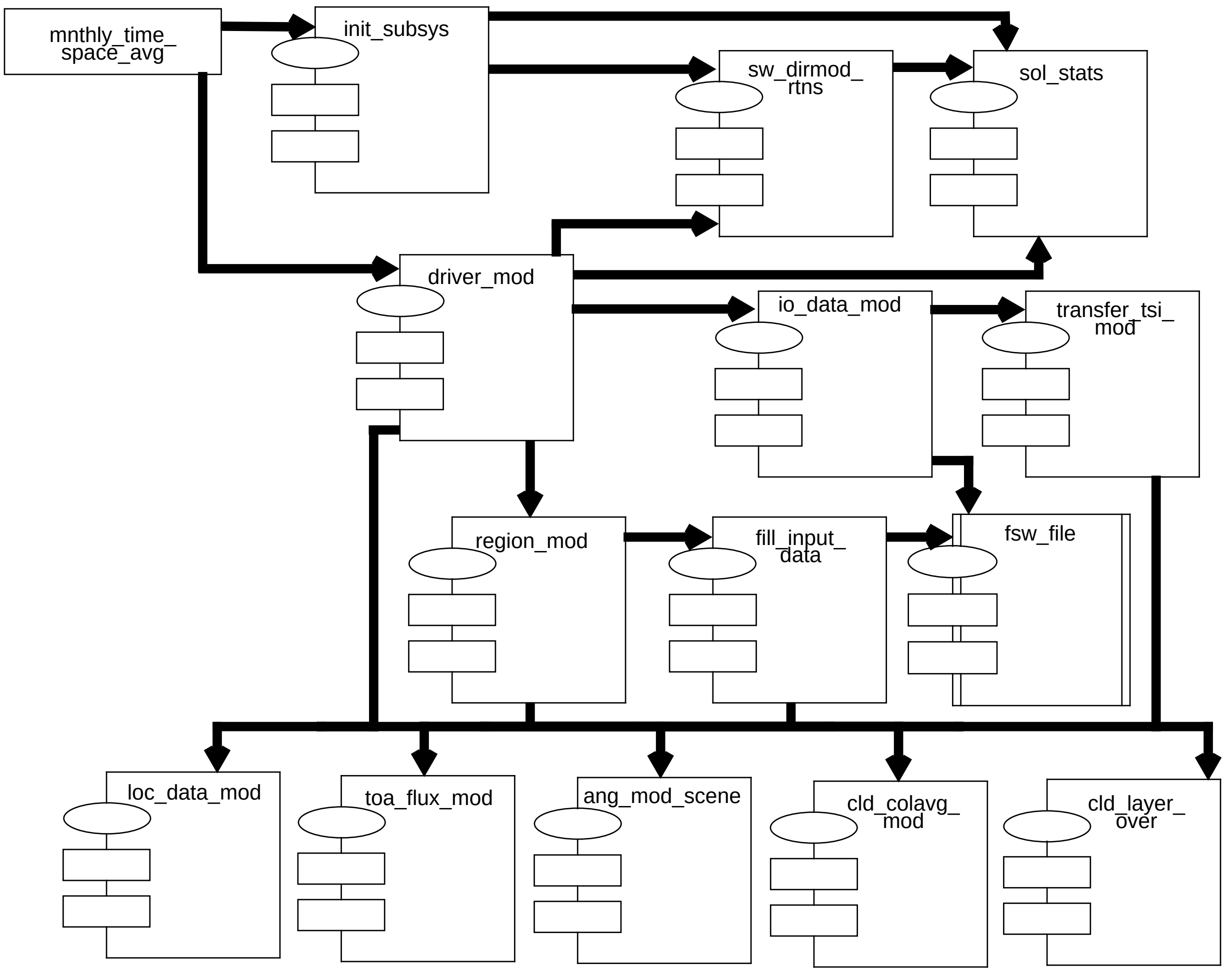

Figure 2-2. Subsystem 7.1 Top-level Context Diagram

2.2 Subsystem 8.0

Figure 2-3 shows the main processing flow of Subsystem 8.0. The data is processed a latitude band, or zone, at a time. For each zone, each region in the zone is processed by reading in the SYN data and then calculating the monthly and monthly hourly regional averages. The zonal averages for the current zone are then calculated and the regional averages written to output. After all zones are processed, the global averages are then calculated. The zonal and global averages are written to output and quality control reports.

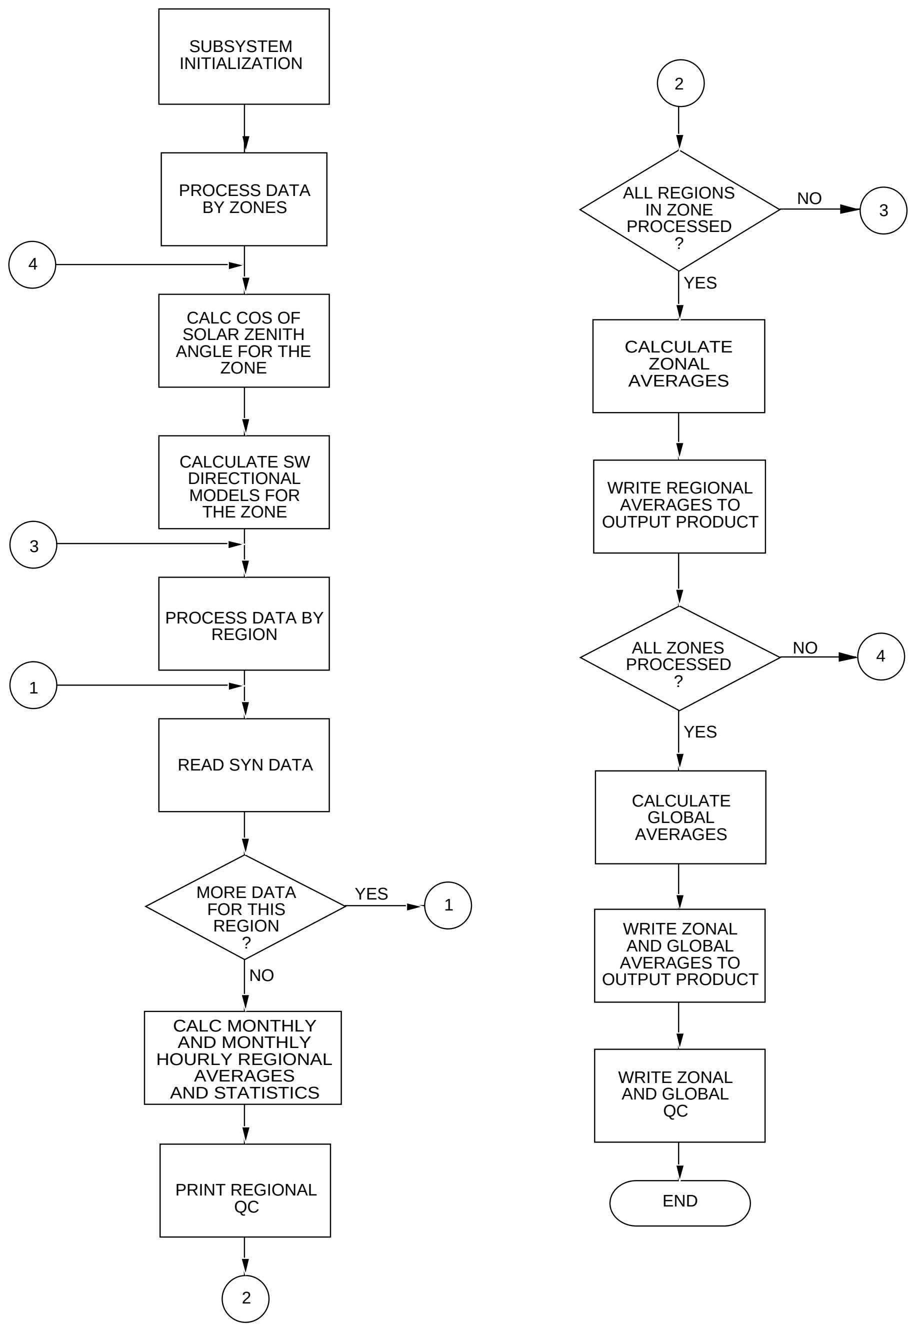

Figure 2-3. Subsystem 8.0 Flow Chart Overview

To introduce the main components of Subsystem 8.0, a context diagram is shown in Figure 2-4. This diagram demonstrates the dependencies between the top level modules that define the data in Subsystem 8.0 and control the main processing, which is as follows : initialize the subsystem; read in the input data, SYN; average by region, zone, and globe; and write the averages to the output products, AVG and ZAVG.

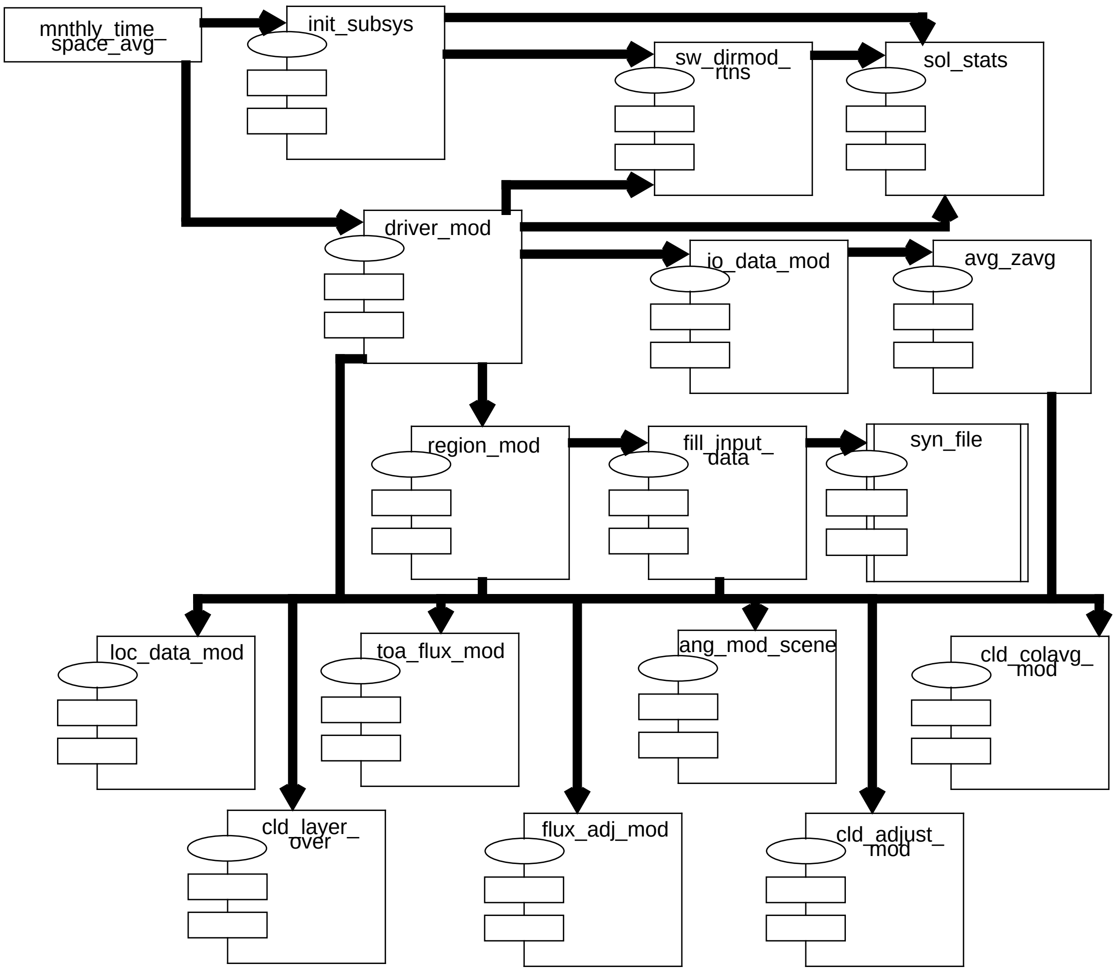

Figure 2-4. Subsystem 8.0 Top-level Context Diagram

2.3 Subsystem 10.0

The main processing flow of Subsystem 10.0 is shown below in Figure 2-5. This flow diagram gives an overview of the three main driver loops: the zonal loop, regional loop, and hour box loop, and the functions performed within these loops.

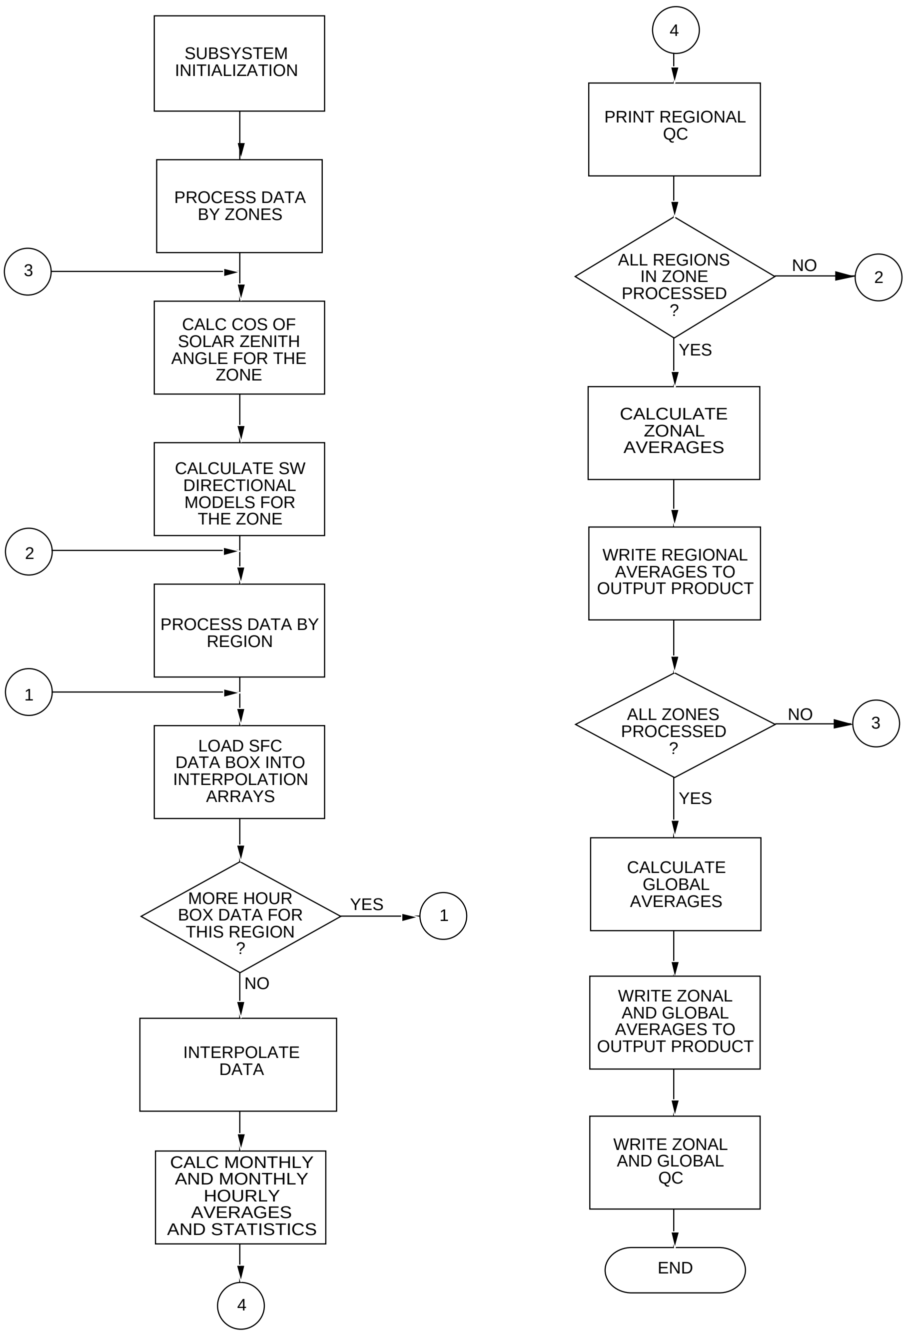

Figure 2-5. Subsystem 10.0 Flow Chart Overview

To introduce the main components of Subsystem 10.0, a context diagram is shown in Figure 2-6. This diagram demonstrates the dependencies between the top level modules that define the data in Subsystem 10.0 and control the main processing, which is as follows : initialize the subsystem; read in the input data, SFC; interpolate the data; average by region, zone, and globe; and write the averages to the output product, SRBAVG.

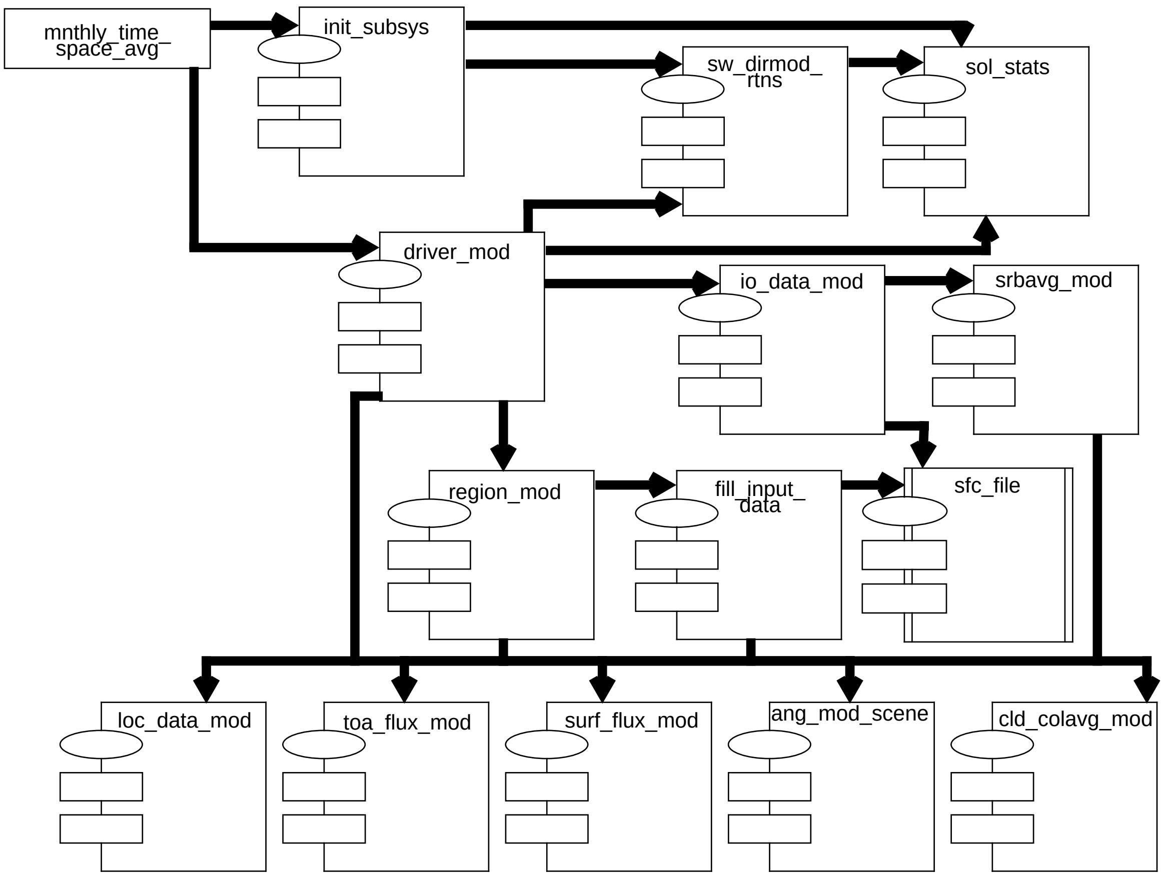

Figure 2-6. Subsystem 10.0 Top-level Context Diagram

2.4 Scenario Diagrams for Subsystems 7.1, 8.0, and 10.0

The next three diagrams are scenario diagrams that demonstrate what is executed at the global level (Figure 2-7), zonal level (Figure 2-8), and regional level (Figure 2-9). The last two scenario diagrams demonstrate which data types are interpolated and regionally averaged for each subsystem (Figure 2-10) and which data types have a zonal and global average for each subsystem (Figure 2-11).

2.4.1 Top-level Senario Diagram

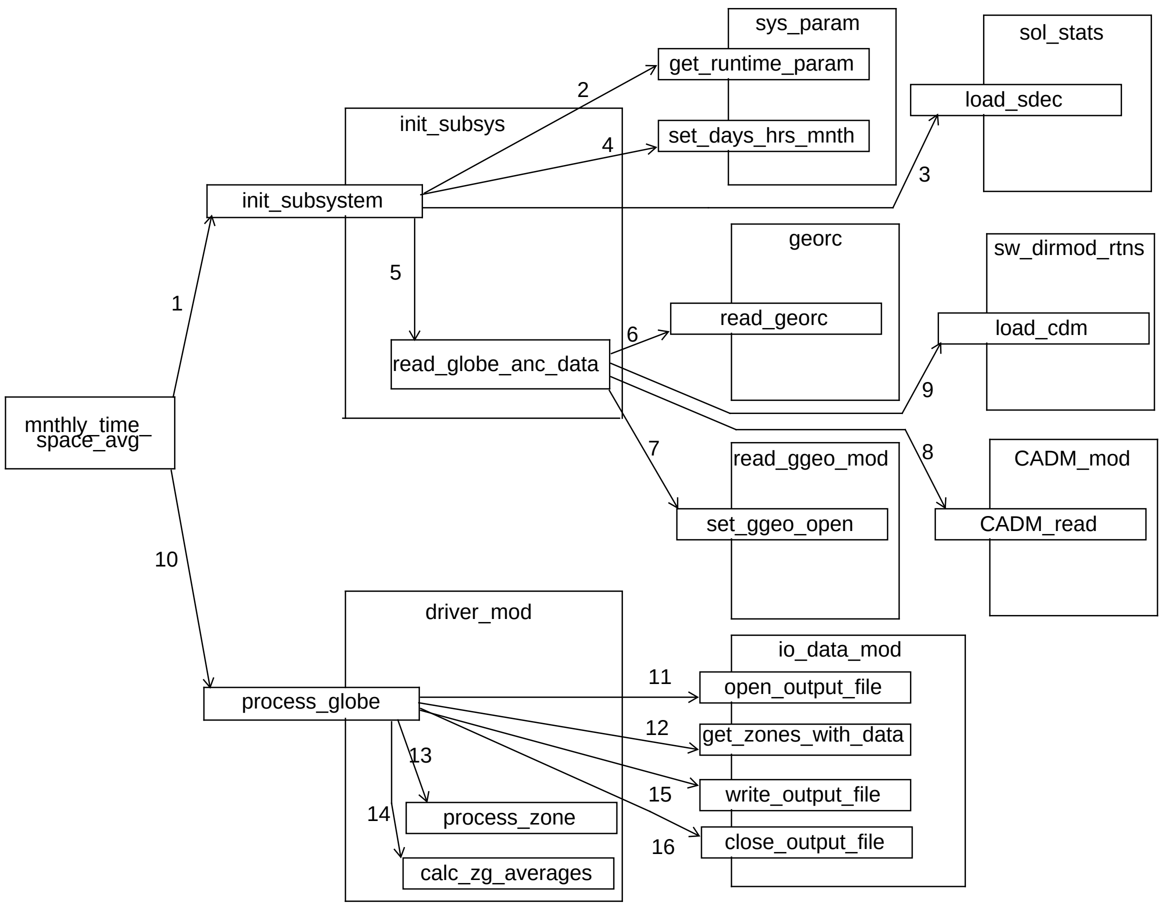

Figure 2-7. Top-level Scenario Diagram

1) The main program, mnthly_time_space_avg calls init_subsystem, which initializes runtime parameters and checks for files existing in the PCF.

2) Get_runtime_param reads in the input parameters from the PCF such as the subsystem number, month, and year being processed.

3) Load_sdec calls the solar_declination program in CERESlib to calculate the distance corrected solar constants and the solar declination angles for the month and stores them in the module sol_stats.

4) Set_days_hrs_mnth sets the number of days and hours in the month depending on the month and the subsystem being processed.

5) Read_globe_anc_data reads in, or initializes data that will be used throughout the entire processing of the globe.

6) Read_georc reads in the geostationary regression coefficients and stores them in the module georc (called only for Subsystems 7.1 and 10.0).

7) Set_ggeo_open opens the file containing the gridded geostationary data (output from Subsystem 11.0), for the globe and the month being processed (called only for Subsystems 7.1 and 10.0)

8) CADM_read opens the CERES anisotropic models (called only for Subsystems 7.1 and 10.0)

9) Load_cdm reads in the CERES Directional Models (CDMs) and stores them in the module sw_dirmod_rtns. (For Release 1 these are the ERBE directional models)

10) Process_globe is the main driver for the processing of CERES data.

11) Open_output_file opens the file that will store the output products : SRBAVG (SS 10), AVG ZAVG (SS 8), and TSI (SS 7.1).

12) Get_zones_with_data gets the number of input zonal files which exist for the month being processed from the Process Control File (PCF).

13) Process_zones processes each region for a zone.

14) Calc_zg_averages calculates the zonal and global averages from the regional averages (called only for Subsystems 8.0 and 10.0).

15) Write_output_file writes the global average to the output product.

16) Close_output_file closes the file containing the output product.

2.4.2 Zonal Level Scenario Diagram

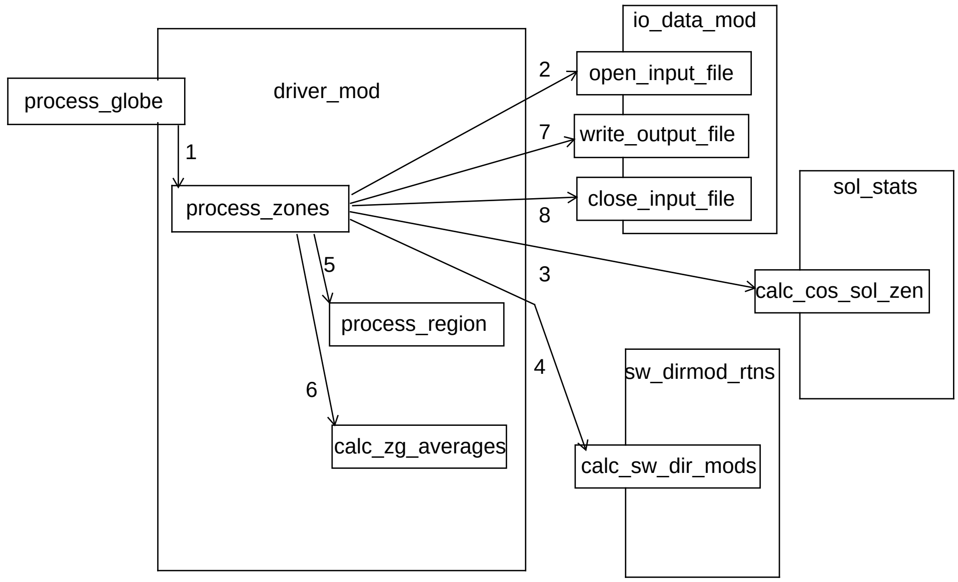

Figure 2-8. Zonal Level Scenario Diagram

1) The main driver subroutine, process_globe, calls process_zones which will do the following steps for every zone in the globe.
2) Open_input_file is called to open the zonal input file, SFC (SS10) and FSW (SS7.1) and the daily input files, SYN (SS8).
3) Calc_cos_sol_zen is called to calculate the cosine of the solar zenith angles for the zone for each hour in the month, and other solar statistics.
4) Calc_sw_dir_mods is called to calculate the Normalized Diurnal Models (NDMs) for the zone for every hour of the month based on the CERES Directional Models (CDM) (called only for Subsystems 7.1 and 10.0).
5) Process_region is called to process each region (interpolate - 7.1 and 10.0, average - 8.0 and 10.0)
6) Calc_zg_averages is called to produce a zonal average from all of the regional averages for the current zone (called for Subsystems 8.0 and 10.0)
7) Write_output_file writes all of the regional averages for the zone to the output file for Subsystems 8.0 and 10.0, or writes an interpolated region at a time to the output product for 7.1, TSI.
8) Once the whole zone has been processed, the input zonal file is closed for Subsystems 7.1 and 10.0. (The input daily files, SYN, for Subsystem 8.0 are kept open until the whole globe has been processed.)

2.4.3 Regional Level Scenario Diagram

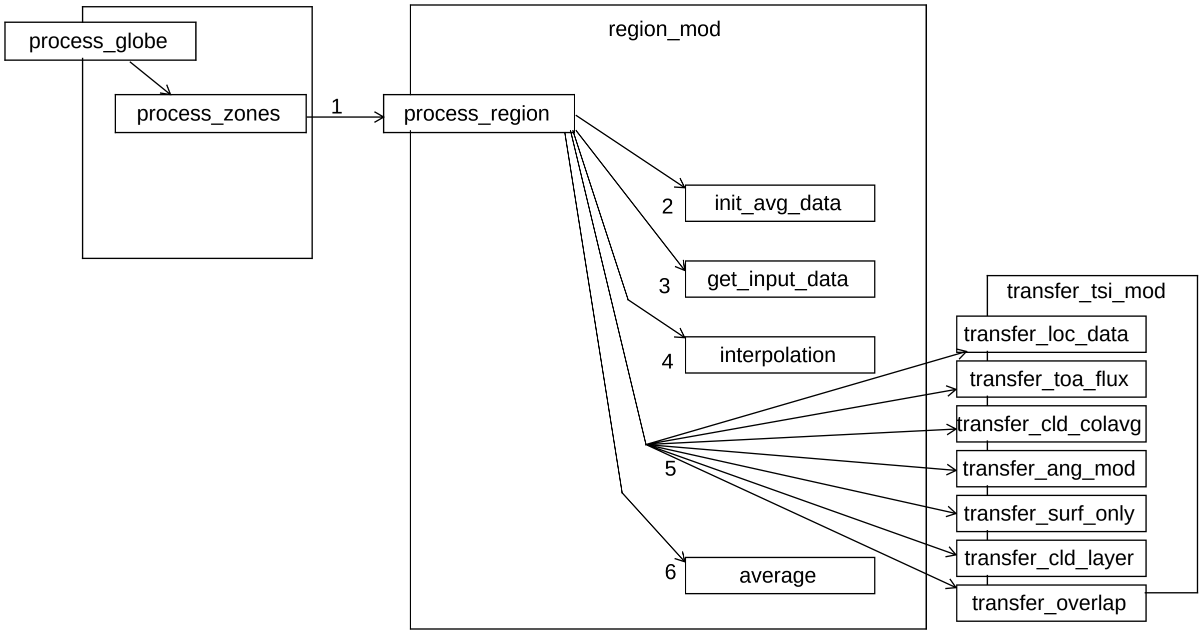

Figure 2-9. Regional Level Scenario Diagram

1) Process_zones calls process_region for each region contained in the zone.

2) Process_region first calls init_avg_data if Subsystems 8.0 or 10.0 are running so that the regional averages for every parameter are set to default values.

3) Get_input_data is called to fill in the input arrays for each parameter from the appropriate input product. This routine calls subroutines from the fill_input_data module.

4) Interpolation is called for Subsystems 7.1 and 10.0, and interpolates each parameter using various algorithms.

5) If Subsystem 7.1 is being run, the transfer routines are called to store the interpolated parameters into the structures defined for the output product for 7.1, TSI.

6) Average is called for Subsystems 8.0 and 10.0, to calculate regional averages of the interpolated data.

2.4.4 Interpolation-Averaging Scenario Diagram

The interpolation-averaging scenario diagram (Figure 2-10) shows which data structures are interpolated and averaged for each region and Subsystem. Note that the interpolation subroutines are called only for Subsystems 7.1 and 10.0, while the averaging subroutines are called only for Subsystems 8.0 and 10.0. The steps used for Subsystem 7.1 are 1-5, 7-11. The steps for Subsystem 8.0 are 12-16, 18-27. The steps for Subsystem 10.0 are 1-8, 12-19.

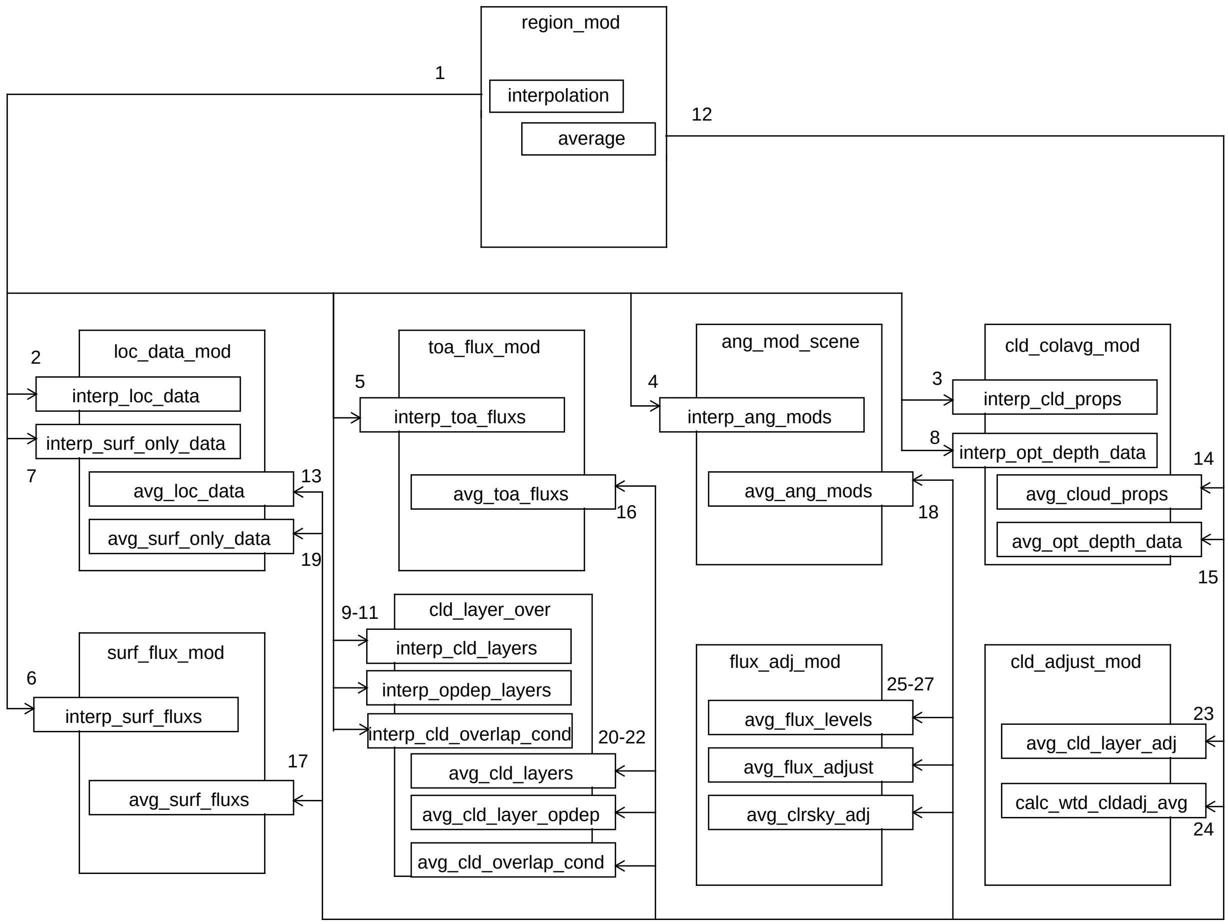

Figure 2-10. Interpolate-Average Scenario Diagram

1) The interpolation subroutine is called for Subsystems 7.1 and 10.0.

2) Interp_loc_data interpolates the location data for Subsystems 7.1 and 10.0.

3) Interp_cld_props interpolates the weighted column averaged cloud properties for Subsystem 7.1 and 10.0.

4) Interp_ang_mods interpolates the angular model scene types for Subsystems 7.1 and 10.0.

5) Interp_toa_fluxs interpolates the top-of-atmosphere fluxes for Subsystems 7.1 and 10.0.

6) Interp_surf_fluxs interpolates the surface fluxes for Subsystem 10.0.
7) Interp_surf_only_data interpolates the surface only data for Subsystems 7.1 and 10.0.
8) Interp_opt_depth interpolates the weighted column averages cloud optical depth data for Subsystems 7.1 and 10.0.
9) Interp_cld_layers interpolates the cloud properties at the four cloud layers for Subsystem 7.1.
10) Interp_opdep_layers interpolates the cloud optical depth data at the four cloud layers for Subsystem 7.1.
11) Interp_cld_overlap_cond interpolates the cloud overlap conditions for Subsystem 7.1.
12) The average subroutine is called for Subsystems 8.0 and 10.0.
13) Avg_loc_data averages the location data for Subsystems 8.0 and 10.0.
14) Avg_cld_props averages the weighted column averaged cloud properties for Subsystems 8.0 and 10.0.
15) Avg_opt_depth_data averages the weighted column averaged cloud optical depth data for Subsystems 8.0 and 10.0.
16) Avg_toa_fluxs averages the top-of-atmosphere fluxes for Subsystems 8.0 and 10.0.
17) Avg_surf_fluxs averages the surface fluxes for Subsystem 10.0.
18) Avg_ang_mods averages the angular model scene type parameters for Subsystems 8.0 and 10.0.
19) Avg_surf_only_data averages the surface only data parameters for Subsystems 8.0 and 10.0.
20) Avg_cld_layers averages the cloud properties at the four cloud layers for Subsystem 8.0.
21) Avg_cld_layer_opdep averages the cloud optical depth data at the four cloud layers for Subsystem 8.0.
22) Avg_cld_overlap_cond averages the cloud overlap conditions for Subsystem 8.0.
23) Avg_cld_layer_adj averages the cloud adjustment parameters at the four cloud layers for Subsystem 8.0.
24) Calc_wtd_cldadj_avg calculates the weighted column averaged cloud adjustment parameters for Subsystem 8.0.
25) Avg_flux_levels averages the fluxes at levels for Subsystem 8.0.
26) Avg_flux_adjust averages the flux adjustment parameters for Subsystem 8.0.
27) Avg_clrsky_adj averages the clear-sky adjustment parameters for Subsystem 8.0.

2.4.5 Zonal-Global-Averaging Scenario Diagram

The zonal-global-averaging scenario diagram (Figure 2-11) shows which data structures for each subsystem are averaged for each zone on the globe. If the input parameters are regional averages, then the subroutines return a zonal average; if the input parameters are zonal averages, the subroutines return a global average. Note that Subsystem 7.1 does not produce any zonal and global averages. The steps for Subsystem 8.0 are 1-2, 4-15. The steps for Subsystem 10.0 are 1-7.

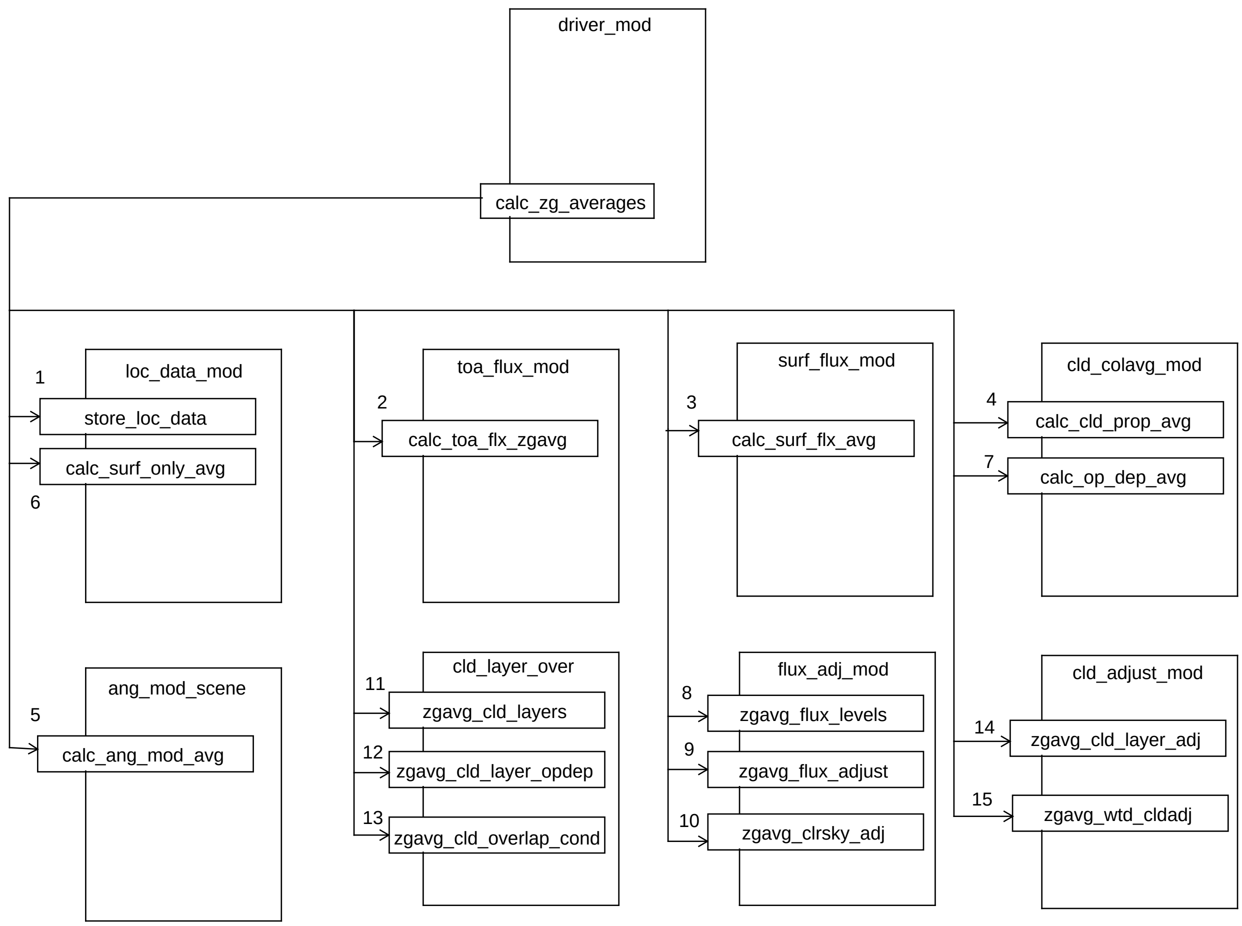

Figure 2-11. Zonal-Global-Average Scenario Diagram

1) Calc_zg_averages calls store_loc_data, which produces either a zonal or global average of the location data.

2) Calc_toa_flx_zgavg calculates the zonal or global average of the top-of-atmosphere fluxes for Subsystems 8.0 and 10.0.

3) Calc_surf_flx_avg calculates the zonal or global average of the surface fluxes for Subsystem 10.0.

4) Calc_cld_prop_avg calculates the zonal or global average of the weighted column averaged cloud properties for Subsystems 8.0 and 10.0.

5) Calc_ang_mod_avg calculates the zonal or global average of the angular model scene type parameters for Subsystems 8.0 and 10.0.

6) Calc_surf_only_avg calculates the zonal or global average of the surface-only data for Subsystems 8.0 and 10.0.

7) Calc_op_dep_avg calculates the weighted column averaged cloud optical depth data for Subsystems 8.0 and 10.0.

8) Zgavg_flux_levels calculates the zonal or global average of the fluxes at levels for Subsystem 8.0.

9) Zgavg_flux_adjust calculates the zonal or global average of the flux adjustment parameters for Subsystem 8.0.

10) Zgavg_clrsky_adj calculates the zonal or global average of the clear-sky adjustment parameters for Subsystem 8.0.

11) Zgavg_cld_layers calculates the zonal or global average of the cloud properties at the four layers for Subsystem 8.0.

12) Zgavg_cld_layer_opdep calculates the zonal or global average of the cloud optical depth values at the four cloud layers for Subsystem 8.0.

13) Zgavg_cld_overlap_cond calculates the zonal or global average of the cloud overlap conditions for Subsystem 8.0.

14) Zgavg_cld_layer_adj calculates the zonal or global average of the cloud adjustment parameters at the four cloud layers for Subsystem 8.0.

15) Zgavg_wtd_cldadj calculates the zonal or global average of the weighted column averaged cloud adjustment parameters for Subsystem 8.0.

References

1. Clouds and the Earth's Radiant Energy System (CERES), Algorithm Theoretical Basis Document, Time Interpolation and Synoptic Flux Computation for Single and Multiple Satellites, Release 1.1, April 1994.
2. Clouds and the Earth's Radiant Energy System (CERES), Algorithm Theoretical Basis Document, ERBE-like Averaging to Monthly TOA and Surface Fluxes, Release 1.1, April 1994.
3. Clouds and the Earth's Radiant Energy System (CERES), Algorithm Theoretical Basis Document, Cloud Optical Property Retrieval, Release 1.1, April 1994.
4. Clouds and the Earth’s Radiant Energy System (CERES), Algorithm Theoretical Basis Document, CERES Inversion to Instantaneous TOA Fluxes, Release 1.1, April 1994.
5. Time Interpolation for Single and Multiple Satellites (Subsystem 7.1), CERES Data Management System Software Requirements Document, Release 1, NASA Langley Research Center, August 1994.
6. Monthly Regional, Zonal, and Global Radiation Fluxes and Cloud Properties, CERES Data Management System Software Requirements Document, Release 1, NASA Langley Research Center, August 1994.
7. Monthly Regional TOA and SRB Averages (Subsystem 10.0), CERES Data Management System Software Requirements Document, Release 1, NASA Langley Research Center, August 1994.
8. Clouds and the Earth’s Radiant Energy System (CERES), Algorithm Theoretical Basis Document, ERBE-Like Inversion to Instantaneous TOA and Surface Fluxes (Subsystem 2.0), Release 1.1, April 1994.
9. Clouds and the Earth’s Radiant Energy System (CERES), Algorithm Theoretical Basis Document, Monthly Regional TOA and Surface Radiation Budget (Subsystem 10.0), Release 1.1, April 1994.
10. CERES Reference Grid, CERES Software Bulletin 95-03, NASA Langley Research Center, May 1995.
11. SDP Toolkit Users Guide for the ECS Project, Hughes Applied Information Systems, Version 1 Final, May 1994.

Appendix A

Abbreviations and Acronyms

Appendix A
Abbreviations and Acronyms

AVG	Monthly Regional Radiative Fluxes and Clouds
CDM	CERES Directional Models
CERES	Clouds and Earth Radiant Energy System
CERESlib	CERES library
CPU	Central Processing Unit
DAAC	Distributed Active Archive Center
deg	degree
EOS	Earth Observing System
EOS-AM	EOS Morning Crossing Mission
EOSDIS	Earth Observing System Data and Information System
EOS-PM	EOS Afternoon Crossing Mission
ERBE	Earth Radiation Budget Experiment
ERBS	Earth Radiation Budget Satellite
F90	FORTRAN90
FSW	Hourly Gridded Single Satellite Fluxes and Clouds
GEORC	GEO Regression Coefficients
GGEO	Grid ISCCP Geostationary Radiances
GMT	Greenwich mean time
HDF	Heirarchical Data Format
IO	Input/Output
ISCCP	International Satellite Cloud Climatology Project
LW	Longwave
NASA	National Aeronautics and Space Administration
NDM	Normalized Diurnal Model
NOAA	National Oceanic and Atmospheric Administration
PCF	Process Control File
PDPS	Planning and Data Production System
SARB	Surface and Atmospheric Radiative Budget
SDP	Science Data Production
SFC	Hourly Gridded Single Satellite TOA and Surface Fluxes
SRB	Surface Radiation Budget
SRBAVG	Monthly and Regional TOA and SRB Averages
SW	Shortwave

SYN	Synoptic Radiative Fluxes and Clouds
TOA	Top-of-the-Atmosphere
TRMM	Tropical Rainfall Measuring Mission
TSI	Time Space Interpolate
ZAVG	Monthly Zonal and Global Radiative Fluxes and Clouds

Appendix B

SYN Day File Design

Appendix B
SYN Day File Design

The SYN input product is divided into a maximum of 31 day files, depending on the number of days in the month. The following Table gives an example of an SYN day file. Each day file contains data for all the regions on the globe for all of the synoptic hours in that day. The synoptic hours in a day are specified as every third hour : 1, 4, 7, ... , 19, 22. The data is written to file in the following: all the regions for hour 1 is first written, then all the regions for hour 4, etc. The records are written in this order because Subsystem 7.2 writes SYN as its output product. Subsystem 7.2 processes one hour at a time for every region, therefore it writes its output product in the order that it processes (refer to Table B-1).

Table B-1. SYN Day File Design

Record #	Region	Hour	Param 1	Param N
1	1	1	...	...
2	2	1	...	...
...	...	...	...	...
26410	26410	1	...	...
26411	1	4	...	...
...	...	...	...	...
52820	26410	4	...	...
...	...	...	...	...
158461	1	19	...	...
...	...	...	...	...
184870	26410	19	...	...
184871	1	22	...	...
...	...	...	...	...
211820	26410	22	...	...

Subsystem 8.0 processed differently from Subsystem 7.2. Subsystem 8.0 needs all of the synoptic hours for a region at one time in order to process. Therefore, in order for Subsystem 8.0 to read the SYN day files, a secondary index file is needed. The SYN secondary index file is used to locate all 248 synoptic hours for a region, as this data is spread over 31 files. The secondary index file contains a record for each region, therefore, there are 26,410 records. Each record contains the record number where each hour for the region can be found on a SYN day file. Hours 1 through 8 of Day 1 are retrieved from SYN Day File 1, and so on.

Table B-2. SYN Secondary Index File Design

Index Record	**Region**	**Hour1 Day1**	**Hour2 Day1**	**Hour... Day1**	**Hour8 Day1**	**Hour... Day...**	**Hour7 Day31**	**Hour8 Day31**
1	1	rec=1	26411	...	184871	...	158461	184871
2	2	2	26412	...	184872	...	158462	184872
3	3	3	26413	...	184873	...	158463	184873
4	4	4	26414	...	184874	...	158464	184874
5	5	5	26415	...	184875	...	158465	184875
6	6	6	26416	...	184876	...	158466	184876
...	...	...	...	...	...	...	...	...
24609	24609	24609	52819	...	211819	...	184869	211819
24610	24610	24610	52820	...	211820	...	184870	211820

Appendix C

TSI File Design

Appendix C
TSI File Design

A TSI record contains data for one region and one hour. Table C-1 shows the regions and hours that are found in each of the 40 files. One TSI file contains a total of 5282 regions and one of the synoptic hours per day for the whole month. The Hour number in the first column in the table refers to the synoptic hour for a day. Therefore, File 1 contains all the first hours of every day of the month for regions 1 - 5282, File 2 contains all the second hours of every day of the month for regions 1 - 5282, etc.

Table C-1. TSI File Layout

Hour for each day	Regions 1 - 5282	Regions 5283 - 10564	Regions 10565 - 15846	Regions 15847 - 21128	Regions 21129 - 26410
1	File 1	File 9	File 17	File 25	File 33
4	File 2	File 10	File 18	File 26	File 34
7	File 3	File 11	File 19	File 27	File 35
10	File 4	File 12	File 20	File 28	File 36
13	File 5	File 13	File 21	File 29	File 37
16	File 6	File 14	File 22	File 30	File 38
19	File 7	File 15	File 23	File 31	File 39
22	File 8	File 16	File 24	File 32	File 40

The following table, Table C-2, gives an example of the record structures in a TSI file. TSI records are only written to file, if there were data for that region and hour. For this example TSI file, the first region that had data was region 12. This file is equivalent to File 1 in the above Table C-1. For this example, we are assuming that the TSI file in Table C-2 contains data for all of the hours for the regions 12 - 5282. The other TSI files would look similar to this example, but for different regions.

Table C-2. Example TSI file

Record	**Region**	**Hrbox**	**Param 1**	**Param 2**	**Param n**
1	12	1	...	...	...
2	12	25	...	...	...
3	12	49	...	...	...
...	...	...	...	...	...
31	12	742	...	...	...
32	13	1	...	...	...
33	13	25	...	...	...
...	13		...	...	...
62	13	742	...	...	...
63	14	1			
...	...	...			
163340	5281	1			
163341	5281	25			
...	...	...			
163370	5281	742			
163371	5282	1			
163372	5282	25			
...	...	....			
163401	5282	742			

Subsystem 7.1 generates a secondary index file which contains the information for a user to determine the location of the data for a given hour in the 40 TSI files. Regional information is not needed in the secondary index file because certain regions are always placed in the same files (refer to Table C-3. In the TSI secondary index file, default values will be written for any regions not observed for an hour. Table C-3 contains a sample TSI secondary index file, that follows the

example in the previous Table C-2. The secondary index file contains a maximum total of 248 records (8 synoptic hours * 31 days = 248); therefore, there is a record for each synoptic hour in the month. Each record in the secondary index file contains an array of record numbers for every region in the globe for that one synoptic hour in the month; therefore, the length of each array is 26, 410 (the number of regions on the globe).

In Table C-2, there were no data for regions 1 - 11, therefore, each synoptic hour array in the secondary index file contains a default value for indices 1 - 11 (contains record information for regions 1 - 11). The following example of a TSI secondary index file in Table C-3, assumes that there are data for every hour of every region ranging from regions 12 - 26,410. This table demonstrates two important points. First, in the column for region 12, the hours for the first day of region 12 are all on record 1, but on 8 different files. Likewise, all of the hours for the second day are on record 2, but again on 8 different files. The second point is that the record numbers for region 5,282 and region 26,410 are different. The record numbers are different because of our example. Remember that there were no data for any of the hours for regions 1 - 12. Hours that do not have data are not written to the TSI file. Therefore, files 1- 8 (refer to Table C-1), are smaller than the rest of the TSI files because they only have a total of 5,270 regions, while the remainder of the files each contain 5,282 regions. Thus, the last record on the files that contain region 26,410's data (files 33-40) is 163,742, while the last record on the files that contain region 5,282's data (files 1-8) is 163,401.

Table C-3. TSI Secondary Index File

*Hrbox	Sec. Index Rec. No.	Reg=1 TSI Rec. No		Reg=12 TSI Rec. No.	Reg=13 TSI Rec. No.		Reg= 5282		Reg= 26410
1	1	default	...	1	32	...	163,371	...	163,712
4	2	default	...	1	32	...	163,371	...	163,712
7	3	default	...	1	32	...	163,371	...	163,712
10	4	default	...	1	32	...	163,371	...	163,712
13	5	default	...	1	32	...	163,371	...	163,712
16	6	default	...	1	32	...	163,371	...	163,712
19	7	default	...	1	32	...	163,371	...	163,712
22	8	default	...	1	32	...	163,371	...	163,712
25	9	default	...	2	33	...	163,372	...	163,713
28	10	default	...	2	33	...	163,372	...	163,713
	...	...	...	...	...	...	...	...	
739	247	default	...	31	62	...	163,401	...	163,742
742	248	default	...	31	62	...	163,401	...	163,742